

IN THE SPECIFICATION:

Please amend the paragraph on page 1, lines 5 - 12 as shown below.

This application is related to co-pending US Patent Application Serial No. [[____]] 10/015,292 which has been issued as United States Patent 6,804,747 ~~(IBM Docket No. AUS920011020US1)~~, entitled APPARATUS AND METHOD OF REDUCING PHYSICAL STORAGE SYSTEMS NEEDED FOR A VOLUME GROUP TO REMAIN ACTIVE by the inventors herein, filed on even date herewith and assigned to the common assignee of this application.

Please replace the description of Fig. 4 on page 6, line 17 with the following:

Fig. 4 is a conceptual view of ~~an LVM~~ a logical volume manager (LVM).

Please replace the paragraph that starts on page 8, line 14 and ends on page 9, line 8 with the following paragraph:

In the depicted example, server 104 is connected to network 102 along with storage unit 106. In addition, clients 108, 110, and 112 are connected to network 102. These clients 108, 110, and 112 may be, for example, personal computers or network computers. In the depicted

AUS920010862US1

example, server 104 provides data, such as boot files, operating system images, and applications to clients 108, 110 and 112. Clients 108, 110 and 112 are clients to server 104. Network data processing system 100 may include additional servers, clients, and other devices not shown. In the depicted example, network data processing system 100 is the Internet with network 102 representing a worldwide collection of networks and gateways that use the Transmission Control Protocol/Internet Protocol (TCP/IP) TCP/IP suite of protocols to communicate with one another. At the heart of the Internet is a backbone of high-speed data communication lines between major nodes or host computers[[host

computers]], consisting of thousands of commercial, government, educational and other computer systems that route data and messages. Of course, network data processing system 100 also may be implemented as a number of different types of networks, such as for example, an intranet, a local area network (LAN), or a wide area network (WAN). Fig. 1 is intended as an example, and not as an architectural limitation for the present invention.

Please replace the paragraph that starts on page 9, line 9 and ends on line 21 with the following paragraph:

Referring to Fig. 2, a block diagram of a data processing system that may be implemented as a server, such as server 104 in Fig. 1, is depicted in accordance with a preferred embodiment of the present invention. Data processing system 200 may be a symmetric multiprocessor (SMP) system including a plurality of processors 202 and 204 connected to system bus 206. Alternatively, a single processor system may be employed. Also connected to system bus 206 is memory controller/cache 208, which provides an interface to local memory 209. Input/output (I/O) I/O bus bridge 210 is connected to system bus 206 and provides an interface to I/O bus 212. Memory controller/cache 208 and I/O bus bridge 210 may be integrated as depicted.

Please replace the paragraph that starts on page 10, line 20 and ends on page 11, line 14 with the following paragraph:

With reference now to Fig. 3, a block diagram illustrating a data processing system is depicted in which the present invention may be implemented. Data processing system 300 is an example of a client computer. Data processing system 300 employs a peripheral component interconnect (PCI) local bus architecture. Although the depicted example employs a PCI bus, other bus architectures such as Accelerated Graphics Port (AGP) and Industry Standard Architecture (ISA) may be used. Processor 302 and main memory 304 are connected to PCI local bus 306 through PCI bridge 308. PCI bridge 308 also may include an integrated memory controller and cache memory for processor 302. Additional connections to PCI local bus 306 may be made through direct component interconnection or through add-in boards. In the depicted example, local area network (LAN) adapter 310, small computer system interface (SCSI) SCSI host bus adapter 312, and expansion bus interface 314 are connected to PCI local bus 306 by direct component connection. In contrast, audio adapter 316, graphics adapter 318, and audio/video adapter 319 are connected to PCI local bus 306 by add-in boards inserted into expansion slots. Expansion bus interface 314 provides a connection for a keyboard and mouse adapter 320, modem 322, and additional memory 324. Small computer system interface (SCSI) host bus adapter 312 provides a connection for hard

AUS920010862US1

disk drive 326, tape drive 328, and compact disk-read-only memory (CD-ROM) ~~CD-ROM~~ drive 330. Typical PCI local bus implementations will support three or four PCI expansion slots or add-in connectors.

Please replace the paragraph that starts on page 12, line 29 and ends on page 13, line 7 with the following paragraph:

To better understand the invention, a more detailed explanation of a logical volume manager (LVM) ~~the LVM~~ is needed. The LVM interacts with application programs and the physical storage devices as shown in Fig. 4. In Fig. 4 three layers are depicted, an application layer 400, a logical layer 410 and a physical layer 420 each having one or more devices. It should be noted that the devices shown in the three layers are not all inclusive. There may be more devices in use in each of the application layer 400 ~~412~~, the logical layer 410 and the physical layer 420 ~~430~~. Thus, the devices in Fig. 4 should be taken only as an example of devices that may be used.

Please replace the paragraph that starts on page 13, line 8 and ends on line 19 with the following paragraph:

The logical layer 410, for all intent and purposes ~~purpose~~, is the LVM. The LVM may be regarded as being made

AUS920010862US1

up of a set of operating system commands, library subroutines or other tools that allow a user to establish and control logical volume storage. The LVM controls physical storage system resources by mapping data between a simple and flexible logical view of storage space and the actual physical storage system. The LVM does this by using a layer of device driver code that runs above traditional device drivers. This logical view of the disk storage is provided to application programs and is independent of the underlying physical disk structure.

Please replace the paragraph that starts on page 13, line 20 and ends on page 14, line 6 with the following paragraph:

The logical layer 410 contains a logical volume 412 that interacts with logical volume device driver 414. A device driver, as is well known in the art, acts as a translator between a device and programs that use the device. That is, the device driver accepts generic commands from programs and translates them into specialized commands for the device. In this case, the logical volume device driver 414 translates commands from an application program that may be executing on the computer system for device driver 430. Thus, when an application program sends commands to file system manager 402 to store or retrieve data from logical volume 412, the file system manager 402 informs the logical volume ~~manager~~ 412 of the application

AUS920010862US1

program's wish. The logical volume ~~manager~~ 412 then conveys the wish to the logical volume device driver 414. The logical volume device driver 414 then consults the appropriate map and instructs the device driver 430 which ones of physical storage systems 422, 424, 426 and 428 to use for the data.

Please replace the paragraph that starts on page 15, line 4 and ends on line 14 with the following paragraph:

Fig. 5 illustrates a first example of a mirroring map that may be used with the present invention. In this example, three physical storage systems (PSS) are used. The physical storage systems are PSS-1 500, PSS-2 510 and PSS-3 520. The mirroring map of Fig. 5 may be named, for example, vgset₁. The three physical storage systems are divided into partitions ~~(see partitions 502, 504 and 506 of PSS-1 500, partitions 512, 514 and 516 of PSS-2 510 and partitions 522, 524 and 526 of PSS-3 520).~~ Mirrored data is stored in the three physical storage systems as illustrated by data A in partitions 502, 512 and 522 of PSS-1 500, PSS-2 510 and PSS-3 520.

Please replace the paragraph that starts on page 15, line 15 and ends on line 26 with the following paragraph:

Fig. 6 illustrates another mirroring map that may be used with the present invention. The mirroring map, as shown in Fig. 6 may be named, for example, vgset₂. In Fig. 6, the physical storage systems are arranged in pairs. Each physical storage system of a pair contains a part of data A. For example, PSS-1 600 is paired off with PSS-4 610. Likewise, PSS-2 620 is paired off with PSS-5 630 and PSS-3 640 is paired off with PSS-6 650. PSS-1 600 contains the first half of data A and PSS-4 610 contains the second half of data A. The data in pairs PSS-2/PSS-5 and PSS-3/PSS-6 mirrors the data in pair PSS-1/PSS-4. This manner of storing data is called striping as a stripe of the data is stored in each physical storage system of a pair.

Please replace the paragraph that starts on page 15, line 27 and ends on line 31 with the following paragraph:

Obviously, both Fig. 5 and Fig. 6 illustrate examples of mirroring maps that may be used with the present invention. They, by no means, are all inclusive. ~~In any case, the mirroring map of Fig. 5 may be named, for example, vgset₁ and that of Fig. 6 vgset₂.~~

Please replace the paragraph that starts on page 17, line 1 and ends on line 15 with the following paragraph:

In the example above, the format for the associate command may be: "associate logical volume₁ vgset₂ 1". ~~Using~~ This command, in essence, instructs the LVM to use mirroring map vgset₂ to store the data of logical volume₁. Consequently, the data will be stored in physical storage systems PSS-1, PSS-2, PSS-3, PSS-4, PSS-5 and PSS-6 as shown in Fig. 7. The number "1" after vgset₂ means that the first partition of the storage systems should be used. Specifying which partition to use is not necessary since without this specification the LVM will use the next available partitions in the physical storage systems. Indeed, because of the high likelihood for errors associated with this option (i.e., forcing the LVM to use a partition that may already be used for another piece of data) only very sophisticated administrators should ever use ~~the~~ this option.

Please replace the paragraph that starts on page 17, line 16 and ends on line 24 with the following paragraph:

Fig. 8 is a flow chart of a process that may be used to create a mirroring map. The process starts when the create command is entered into the computer system (step 800). A check is made to determine whether a name is provided for the mirroring map being created (step 805). If not, an error such as " name is needed" may be generated (steps ~~805 and~~ 810). If a name is provided, then the user

or administrator will be allowed to define the map. To do so, a display of a physical storage system may be provided.

Please replace the paragraph that starts on page 17, line 25 and ends on page 18, line 3 with the following paragraph:

Using copies (by cutting and pasting, for example) of the provided physical storage system, the administrator may make any mirroring map desired. After defining a map, the administrator has to specify which physical storage systems to be used by providing the names ~~name~~ (each physical storage system has a name) of the physical storage systems (e.g., PSS-1, PSS-2 etc.). Thus, the administrator may use two or more physical storage systems for one or more mirrored copies of the data, and only one physical storage system for the other mirror or mirrors of the data or any combination thereof.

Please replace the paragraph that starts on page 18, line 19 and ends on line 27 with the following paragraph:

If the administrator wants to stripe the data in three stripes within only one physical storage system, the administrator may enter "A/3" into three different partitions of that storage system. Thus, the administrator may instruct the LVM to mirror the data into the physical

AUS920010862US1

storage systems in any combination the Administrator may choose (step 815). When done, the name of the mirroring map and the mirroring map itself are stored in the LVM for future use (~~steps 805, 815 and 820~~).

Please replace the paragraph that starts on page 18, line 28 and ends on page 19, line 4 with the following paragraph:

Fig. 9 is a flow diagram of a process that may be used to display the mirroring maps. The process starts as soon as the display command is entered into the computer system (step 900). A check is then made to determine whether a mirroring map name is provided (step 905). If so the mirroring map associated with the name entered and the name of the mirroring map will all be displayed (~~steps 905 and 910~~). If not, all the mirroring maps stored in the LVM and their names will be displayed (~~steps 905 and 915~~).

Please replace the paragraph that starts on page 19, line 5 and ends on line 28 with the following paragraph:

Fig 10 is a flow diagram of a process that may be used to associate a logical volume with a mirroring map. The process starts when the associate command is entered into the computer system (step 1000). A first check is then made to determine whether a name of a logical volume is

AUS920010862US1

provided with the command (step 1005). If not, an error such as "a name of a logical volume is needed" may be generated (~~steps 10005 and~~ 1010). If a name of a logical volume is provided with the command or after the name is entered, a second check is made to determine whether the name of the mirroring map is provided with the command (step 1015). If not, then an error such as "please enter a name of a mirroring map" may be generated (~~steps 10005, 1010, 1015 and~~ 1020). If the name of a mirroring map is provided with the command or after entering the name of the mirroring map, a third check is made to determine whether the mirroring map exists (step 1025). If not, an error such as "mirroring map does not exist" may be generated (~~steps 1025 and~~ 1030). If the mirroring map does exist, then the name of the logical volume provided is associated with the mirroring map and stored in the LVM for future reference (step 1035). That is, each time a piece of data is stored in the logical volume, the physical storage systems in the mirroring map will be used to store the data (~~steps 1025, 1030 and~~ 1040).

Please replace the paragraph that starts on page 20, line 8 and ends on line 19 with the following paragraph:

When an attempt to write metadata into a physical storage system is not successful, the physical storage system is marked as unavailable. The present policy is that when half or more of the physical storage systems

AUS920010862US1

present that make up a volume group is marked unavailable, the volume group should not be used or remain active. This policy is used to ensure that when the computer system is turned on or is reset, if the majority of the physical storage systems in a volume group is available, the volume group may be used as metadata in at least one of the physical storage systems will be valid. The valid metadata may be copied into the other physical storage systems.

Please replace the paragraph that starts on page 21, line 18 and ends on line 28 with the following paragraph:

Fig. 11 is a flow diagram of a process that may be used to ascertain ~~that~~ whether there is a requisite number of physical storage systems available for a volume group to remain active. The process starts when a computer system is turned on or reset (step 1100). A check is continuously being made to determine whether all the physical storage systems in use are available (step 1105). If not, a check will be made to determine whether at least one physical storage system from each of mirroring set is available (step 1115). If so, the volume group will remain active (step 1120). If not, the volume group should not be ~~used~~ used or remain active anymore (~~steps 1105— 1125~~).

Please replace the paragraph that starts on page 21, line 29 and ends on page 22, line 10 with the following paragraph:

Fig. 12 is a flow diagram of a process that may be used to determine whether a volume group has a requisite number of physical storage systems available when a computer system is turned on or reset. Again the process starts when the computer system is turned on or reset (step 1200). A check is made to determine whether all the physical storage systems in the volume group are available (step 1205). If not, it is then determined whether an entire mirror set is available (step 1215). If not, the volume group may not activate (~~steps 1205—~~ 1220). If there is an entire mirror set available, it will then be determined whether at least one physical storage system out of each of the remaining mirroring sets is available (step 1225). If not, the volume group may not activate (step 1230). If so, the volume group may activate (~~steps 1225—~~ 1235).

In Fig. 10, please insert -- ASSOCIATE COMMAND ENTERED -- after "START" in item 1000;

please insert -- ERROR -- before "A NAME" in box 1010;

please insert -- ERROR -- before "A NAME" in box 1020;

please insert -- ERROR -- before "MAP" in box 1030.

In Fig. 11, please insert -- SYSTEM ON OR RESET -- after "START" in item 1100;

please delete "OK" in box 1110. Further, delete box 1110 and reference numeral "1110" and replace with arrow that goes from decision box 1105 to arrow pointing to decision box 1105;

please replace "OK" in box 1120 with -- VOLUME GROUP SHOULD REMAIN ACTIVE --.

In Fig. 12, please insert -- SYSTEM ON OR RESET -- after "START" in item 1200;

please delete "OK" in box 1210. Further, delete box 1210 and reference numeral "1210" and replace with arrow that goes from decision box 1205 to arrow pointing to decision box 1205.